

03-28-00

PATENT

A

03/27/00
jc658 U.S. PTO

Attorney's Docket No. 99-412

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

jc530 U.S. PTO
09/536191
03/27/00

Box Patent Application
Assistant Commissioner for Patents
Washington, D. C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of:

Inventor(s): Robert J. Donaghey

For: SYSTEMS AND METHODS FOR COMMUNICATING IN A
PERSONAL AREA NETWORK

Certification Under 37 CFR 1.10

I hereby certify that this New Application Transmittal and the documents referred to as enclosed therein are being deposited with the United States Postal Service on this date 3/27/00 in an envelope as "Express Mail Post Office to Addressee" mailing label EK006159886US addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

Mary E. Anza
(typed or printed name of person mailing paper)


Mary E. Anza
(Signature of person mailing paper)

Attorney's Docket No. 99-412

Enclosed are:

[26] pages of specification and cover sheet
 [11] pages of claims
 [1] page of abstract
 [17] sheets of formal drawings.
 [2] pages of declaration and power of attorney.
 [2] pages of assignment and assignment recordation form
 [2] pages of information disclosure statement
 [1] page of form 1449
 [1] reference
 [1] return postcard

CLAIMS AS FILED			
	NUMBER FILED	NUMBER EXTRA	BASIC FEE \$690.00
TOTAL CLAIMS	43 - 20	23 x	\$18.00
INDEPENDENT CLAIMS	9 - 3	6 x	\$78.00
MULTIPLE DEPENDENT CLAIMS(S)		0 x	\$260.00
TOTAL FILING FEE			\$1572.00

- 
Floyd E. Anderson
Registration No. 33,825
Attorney for Applicant(s)

Year	1950	1951	1952	1953	1954	1955	1956	1957	1958	1959	1960	1961	1962	1963	1964	1965	1966	1967	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100
1950	1951	1952	1953	1954	1955	1956	1957	1958	1959	1960	1961	1962	1963	1964	1965	1966	1967	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	

PATENT
Docket No. 99-412

OF

Robert J. DONAGHEY

FOR

SYSTEMS AND METHODS FOR COMMUNICATING IN A PERSONAL AREA NETWORK

[illegible]

SYSTEMS AND METHODS FOR COMMUNICATING
IN A PERSONAL AREA NETWORK

RELATED APPLICATION

This application is related to U.S. Patent Application, Serial No. _____

5 (Atty. Docket No. 99-425), entitled "PERSONAL AREA NETWORK WITH AUTOMATIC
ATTACHMENT AND DETACHMENT," filed on an even date herewith, and incorporated
herein by reference.

BACKGROUND OF THE INVENTION

A. Field of the Invention

10 The present invention relates to a data network and, more particularly, to a star data
network that facilitates bidirectional wireless data communications between a main processor
unit and a varying number of peripheral units as they become located within the proximity of the
processor unit.

B. Description of Related Art

15 Over the last decade, the size and power consumption of digital electronic devices has
been progressively reduced. For example, personal computers have evolved from lap tops and
notebooks into hand-held or belt-carriable devices commonly referred to as personal digital
assistants (PDAs). One area of carriable devices that has remained troublesome, however, is the
coupling of peripheral devices or sensors to the main processing unit of the PDA. Generally,
20 such coupling is performed through the use of connecting cables. The connecting cables restrict
the handling of a peripheral in such a manner as to lose many of the advantages inherent in the

PDA's small size and light weight. For a sensor, for example, that occasionally comes into contact with the PDA, the use of cables is particularly undesirable.

While some conventional systems have proposed linking a keyboard or a mouse to a main processing unit using infrared or radio frequency (RF) communications, such systems have typically been limited to a single peripheral unit with a dedicated channel of low capacity.

Based on the foregoing, it is desirable to develop a low power data network that provides highly reliable bidirectional data communication between a host or server processor unit and a varying number of peripheral units and/or sensors while avoiding interference from nearby similar systems.

SUMMARY OF THE INVENTION

Systems and methods consistent with the present invention address this need by providing a wireless personal area network that permits a host unit to communicate with peripheral units with minimal interference from neighboring systems.

A system consistent with the present invention includes a hub device connected to several peripheral devices via a network. The hub device generates a token and broadcasts the token on the network. Each of the peripheral devices receives the token broadcast by the hub device, determines whether the token identifies the peripheral device, analyzes the token to determine a size and direction of a current data transfer when the token identifies the peripheral device, and transfers data to or receives data from the hub device according to the determined size and direction of the current data transfer.

In another implementation consistent with the present invention, a memory, used by at least one of a hub device and a peripheral device to transfer data between the hub device and multiple peripheral devices over a wireless network, includes a network interface structure, a link layer transport structure, and a link layer driver structure. The network interface structure determines whether and when to schedule a data transfer. The link layer transport structure receives the data from the network interface structure and arranges for a reliable transmission of the data. The link layer driver structure receives the data from the link layer transport structure and transmits the data under control of the link layer transport structure.

In yet another implementation consistent with the present invention, a communications protocol, used in a network connecting a hub device to at least one peripheral device, includes several frames, including a beacon, at least one token transmission, and at least one data transfer opportunity. The beacon marks a start of the corresponding frame. The token transmission identifies one of the peripheral devices for a data transfer. The data transfer opportunity permits the hub device to communicate a data block with the identified peripheral device.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an embodiment of the invention and, together with the description, explain the invention. In the drawings:

Fig. 1 is a diagram of a personal area network (PAN) in which systems and methods consistent with the present invention may be implemented;

Fig. 2 is a simplified block diagram of the Hub of Fig. 1;

Fig. 3 is a simplified block diagram of a PEA of Fig. 1;

Fig. 4 is a block diagram of a software architecture of a Hub or PEA in an implementation consistent with the present invention;

Fig. 5 is an exemplary diagram of communication processing by the layers of the software architecture of Fig. 4;

Fig. 6 is an exemplary diagram of a data block architecture within the DCL of the Hub and PEA in an implementation consistent with the present invention;

Fig. 7A is a detailed diagram of an exemplary stream usage plan in an implementation consistent with the present invention;

Fig. 7B is a detailed diagram of an exemplary stream usage assignment in an implementation consistent with the present invention;

Fig. 8 is an exemplary diagram of a time division multiple access (TDMA) frame structure in an implementation consistent with the present invention;

Fig. 9A is a detailed diagram of activity within the Hub and PEA according to a TDMA plan consistent with the present invention;

Fig. 9B is a flowchart of the Hub activity of Fig. 9A;

Fig. 9C is a flowchart of the PEA activity of Fig. 9A;

Figs. 10A and 10B are high-level diagrams of states that the Hub and PEA traverse during a data transfer in an implementation consistent with the present invention;

Figs. 11 and 12 are flowcharts of Hub and PEA attachment processing, respectively, consistent with the present invention; and

Fig. 13 is a flowchart of PEA detachment and reattachment processing consistent with the present invention.

DETAILED DESCRIPTION

5 The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims.

10 Systems and methods consistent with the present invention provide a wireless personal area network that permits a host device to communicate with a varying number of peripheral devices with minimal interference from neighboring networks. The host device uses tokens to manage all of the communication in the network.

NETWORK OVERVIEW

15 A Personal Area Network (PAN) is a local network that interconnects computers with devices (*e.g.*, peripherals, sensors, actuators) within their immediate proximity. These devices may be located nearby and may frequently or occasionally come within range and go out of range of the computer. Some devices may be embedded within an infrastructure (*e.g.*, a building or vehicle) so that they can become part of a PAN as needed.

20 A PAN, in an implementation consistent with the present invention, has low power consumption and small size, supports wireless communication without line-of-sight limitations, supports communication among networks of multiple devices (over 100 devices), and tolerates interference from other PAN systems operating within the vicinity. A PAN can also be easily

integrated into a broad range of simple and complex devices, is low in cost, and is capable of being used worldwide.

Fig. 1 is a diagram of a PAN 100 consistent with the present invention. The PAN 100 includes a single Hub device 110 surrounded by multiple Personal Electronic Accessory (PEA) devices 120 configured in a star topology. Other topologies may also be possible. Each device is identified by a Media Access (MAC) address.

The Hub 110 orchestrates all communication in the PAN 100, which consists of communication between the Hub 110 and one or more PEA(s) 120. The Hub 110 manages the timing of the network, allocates available bandwidth among the currently attached PEAs 120 participating in the PAN 100, and supports the attachment, detachment, and reattachment of PEAs 120 to and from the PAN 100.

The Hub 110 may be a stationary device or may reside in some sort of wearable computer, such as a simple pager-like device, that may move from peripheral to peripheral. The Hub 110 could, however, include other devices.

The PEAs 120 may vary dramatically in terms of their complexity. A very simple PEA might include a movement sensor having an accelerometer, an 8-bit microcontroller, and a PAN interface. An intermediate PEA might include a bar code scanner and its microcontroller. More complex PEAs might include PDAs, cellular telephones, or even desktop PCs and workstations. The PEAs may include stationary devices located near the Hub and/or portable devices that move to and away from the Hub.

The Hub 110 and PEAs 120 communicate using multiplexed communication over a predefined set of streams. Logically, a stream is a one-way communications link between one

PEA 120 and its Hub 110. Each stream has a predetermined size and direction. The Hub 110 uses stream numbers to identify communication channels for specific functions (*e.g.*, data and control).

The Hub 110 uses MAC addresses to identify itself and the PEAs 120. The Hub 110 uses its own MAC address to broadcast to all PEAs 120. The Hub 110 might also use MAC addresses to identify virtual PEAs within any one physical PEA 120. The Hub 110 combines a MAC address and a stream number into a token, which it broadcasts to the PEAs 120 to control communication through the network 100. The PEA 120 responds to the Hub 110 if it identifies its own MAC address or the Hub MAC address in the token and if the stream number in the token is active for the MAC address of the PEA 120.

EXEMPLARY HUB DEVICE

Fig. 2 is a simplified block diagram of the Hub 110 of Fig. 1. The Hub 110 may be a battery-powered device that includes Hub host 210, digital control logic 220, radio frequency (RF) transceiver 230, and an antenna 240.

Hub host 210 may include anything from a simple microcontroller to a high performance microprocessor. The digital control logic (DCL) 220 may include a controller that maintains timing and coordinates the operations of the Hub host 210 and the RF transceiver 230. The DCL 220 is specifically designed to minimize power consumption, cost, and size of the Hub 110. Its design centers around a time-division multiple access (TDMA)-based network access protocol that exploits the short range nature of the PAN 100. The Hub host 210 causes the DCL 220 to initialize the network 100, send tokens and messages, and receive messages. Responses from the DCL 220 feed incoming messages to the Hub host 210.

The RF transceiver 230 includes a conventional RF transceiver that transmits and receives information via the antenna 240. The RF transceiver 230 may alternatively include separate transmitter and receiver devices controlled by the DCL 220. The antenna 240 includes a conventional antenna for transmitting and receiving information over the network.

5 While Fig. 2 shows the exemplary Hub 110 as consisting of three separate elements, these elements may be physically implemented in one or more integrated circuits. For example, the Hub host 210 and the DCL 220, the DCL 220 and the RF transceiver 230, or the Hub host 210, the DCL 220, and the RF transceiver 230 may be implemented as a single integrated circuit or separate integrated circuits. Moreover, one skilled in the art will recognize that the Hub 110 may
10 include additional elements that aid in the sending, receiving, and processing of data.

EXEMPLARY PEA DEVICE

Fig. 3 is a simplified block diagram of the PEA 120. The PEA 120 may be a battery-powered device that includes a PEA host 310, DCL 320, RF transceiver 330, and an antenna 340. The PEA host 310 may include a sensor that responds to information from a user, an actuator
15 that provides output to the user, a combination of a sensor and an actuator, or more complex circuitry, as described above.

The DCL 320 may include a controller that coordinates the operations of the PEA host 310 and the RF transceiver 330. The DCL 320 sequences the operations necessary in establishing synchronization with the Hub 110, in data communications, in coupling received
20 information from the RF transceiver 330 to the PEA host 310, and in transmitting data from the PEA host 310 back to the Hub 110 through the RF transceiver 330.

The RF transceiver 330 includes a conventional RF transceiver that transmits and receives information via the antenna 340. The RF transceiver 330 may alternatively include separate transmitter and receiver devices controlled by the DCL 320. The antenna 340 includes a conventional antenna for transmitting and receiving information over the network.

5 While Fig. 3 shows the exemplary PEA 120 as consisting of three separate elements, these elements may be physically implemented in one or more integrated circuits. For example, the PEA host 310 and the DCL 320, the DCL 320 and the RF transceiver 330, or the PEA host 310, the DCL 320, and the RF transceiver 330 may be implemented as a single integrated circuit or separate integrated circuits. Moreover, one skilled in the art will recognize that the PEA 120
10 may include additional elements that aid in the sending, receiving, and processing of data.

EXEMPLARY SOFTWARE ARCHITECTURE

Fig. 4 is an exemplary diagram of a software architecture 400 of the Hub 110 in an implementation consistent with the present invention. The software architecture 400 in the PEA 120 has a similar structure. The software architecture 400 includes several distinct layers, each
15 designed to serve a specific purpose, including: (1) application 410, (2) link layer control (LLC) 420, (3) network interface (NI) 430, (4) link layer transport (LLT) 440, (5) link layer driver (LLD) 450, and (6) DCL hardware 460. The layers have application programming interfaces (APIs) to facilitate communication with lower layers. The LLD 450 is the lowest layer of software. Each layer may communicate with the next higher layer via procedural upcalls that the
20 higher layer registers with the lower layer.

The application 410 may include any application executing on the Hub 110, such as a communication routine. The LLC 420 performs several miscellaneous tasks, such as

initialization, attachment support, bandwidth control, and token planning. The LLC 420 orchestrates device initialization, including the initialization of the other layers in the software architecture 400, upon power-up.

The LLC 420 provides attachment support by providing attachment opportunities for unattached PEAs to attach to the Hub 110 so that they can communicate, providing MAC address assignment, and initializing an NI 430 and the layers below it for communication with a PEA 120. The LLC 420 provides bandwidth control through token planning. Through the use of tokens, the LLC 420 allocates bandwidth to permit one PEA 120 at a time to communicate with the Hub 110.

The NI 430 acts on its own behalf, or for an application 410 layer above it, to deliver data to the LLT 440 beneath it. The LLT 440 provides an ordered, reliable "snippet" (*i.e.*, a data block) delivery service for the NI 430 through the use of encoding (*e.g.*, 16-64 bytes of data plus a cyclic redundancy check (CRC)) and snippet retransmission. The LLT 440 accepts snippets, in order, from the NI 430 and delivers them using encoded status blocks (*e.g.*, up to 2 bytes of status information translated through Forward Error Correction (FEC) into 6 bytes) for acknowledgments (ACKs).

The LLD 450 is the lowest level of software in the software architecture 400. The LLD 450 interacts with the DCL hardware 460. The LLD 450 initializes and updates data transfers via the DCL hardware 460 as it delivers and receives data blocks for the LLT 440, and processes hardware interrupts. The DCL hardware 460 is the hardware driven by the LLD 450.

Fig. 5 is an exemplary diagram of communication processing by the layers of the software architecture 400 of Fig. 4. In Fig. 5, the exemplary communications involve the transmission of

a snippet from one node to another. This example assumes that the sending node is the Hub 110 and the receiving node is a PEA 120. Processing begins with the NI 430 of the Hub 110 deciding to send one or more bytes (but no more than will fit) in a snippet. The NI 430 exports the semantics that only one transaction is required to transmit these bytes to their destination

5 (denoted by "(1)" in the figure). The NI 430 sends a unique identifier for the destination PEA 120 of the snippet to the LLT 440. The LLT 440 maps the PEA identifier to the MAC address assigned to the PEA 120 by the Hub 110.

The LLT 440 transmits the snippet across the network to the receiving device. To accomplish this, the LLT 440 adds header information (to indicate, for example, how many bytes
10 in the snippet are padded bytes) and error checking information to the snippet, and employs reverse-direction status/acknowledgment messages and retransmissions. This is illustrated in Fig. 5 by the bidirectional arrow between the LLT 440 layers marked with "(n + m)." The number n of snippet transmissions and the number m of status transmissions in the reverse direction are mostly a function of the amount of noise in the wireless communication, which may
15 be highly variable. The LLT 440 may also encrypt portions or all of the snippet using known encryption technology.

The LLT 440 uses the LLD 450 to provide a basic block and stream-oriented communications service, isolating the DCL 460 interface from the potentially complex processing required of the LLT 440. The LLT 440 uses multiple stream numbers to differentiate
20 snippet and status blocks so that the LLD 450 need not know which blocks contain what kind of content. The LLD 450 reads and writes the hardware DCL 460 to trigger the transmission and reception of data blocks. The PEA LLT 440, through the PEA LLD 450, instructs the PEA DCL

460 which MAC address or addresses to respond to, and which stream numbers to respond to for each MAC address. The Hub LLT 440, through the Hub LLD 450, instructs the Hub DCL 460 which MAC addresses and stream numbers to combine into tokens and transmit so that the correct PEA 120 will respond. The Hub DCL 460 sends and receives (frequently in a corrupted form) the data blocks across the RF network via the Hub RF transceiver 230 (Fig. 2).

The Hub LLT 440 employs FEC for status, checksums and error checking for snippets, and performs retransmission control for both to ensure that each snippet is delivered reliably to its client (e.g., PEA LLT 440). The PEA LLT 440 delivers snippets in the same order that they were sent by the Hub NI 430 to the PEA NI 430. The PEA NI 430 takes the one or more bytes sent in the snippets and delivers them in order to the higher-level application 410, thereby completing the transmission.

EXEMPLARY DCL DATA BLOCK ARCHITECTURE

Fig. 6 is an exemplary diagram of a data block architecture 600 within the DCL of the Hub 110 and the PEA 120. The data block 600 contains a MAC address 610 designating a receiving or sending PEA 120, a stream number 620 for the communication, and a data buffer 630 which is full when sending and empty when receiving. As will be described later, the MAC address 610 and stream number 620 form the contents of a token 640. When the LLD 450 reads from and writes to the hardware DCL 460, the LLD 450 communicates the MAC address 610 and stream number 620 with the data buffer 630. When a PEA 120 receives a data block, the DCL 460 places the MAC address 610 and stream number 620 contained in the preceding token 640 in the data block 600 to keep track of the different data flows.

EXEMPLARY STREAM ARCHITECTURE

The LLD 450 provides a multi-stream data transfer service for the LLT 440. While the LLT 440 is concerned with data snippets and status/acknowledgements, the LLD 450 is concerned with the size of data blocks and the direction of data transfers to and from the Hub 110.

Fig. 7A is a detailed diagram of an exemplary stream usage plan 700 in an implementation consistent with the present invention. A single stream usage plan may be predefined and used by the Hub 110 and all PEAs 120. The PEA 120 may have a different set of active streams for each MAC address it supports, and only responds to a token that specifies a MAC address of the PEA 120 and a stream that is active for that MAC address. In an implementation consistent with the present invention, every PEA 120 may support one or more active Hub-to-PEA streams associated with the Hub's MAC address.

The stream usage plan 700 includes several streams 710-740, each having a predefined size and data transfer direction. The plan 700 may, of course, have more or fewer entries and may accommodate more than the two data block sizes shown in the figure. In the plan 700, streams 0-2 (710) are used to transmit the contents of small data blocks from the PEA 120 to the Hub 110. Streams 3-7 (720) are used to transmit the contents of larger data blocks from the PEA 120 to the Hub 110. Streams 8-10 (730), on the other hand, are used to transmit the contents of small data blocks from the Hub 110 to the PEA 120. Streams 11-15 (740) are used to transmit the contents of larger data blocks from the Hub 110 to the PEA 120.

To avoid collisions, some of the streams are reserved for PEAs desiring to attach to the network and the rest are reserved for PEAs already attached to the network. With such an

arrangement, a PEA 120 knows whether and what type of communication is scheduled by the Hub 110 based on a combination of the MAC address 610 and the stream number 620.

Fig. 7B is a detailed diagram of an exemplary stream usage assignment by the LLT 440 in an implementation consistent with the present invention. The LLT 440 assigns different streams to different communication purposes, reserving the streams with small block size for status, and using the streams with larger block size for snippets. For example, the LLT 440 may use four streams (4-7 and 12-15) for the transmission of snippets in each direction, two for odd parity snippets and two for even parity snippets. In other implementations consistent with the present invention, the LLT 440 uses different numbers of streams of each parity and direction.

The use of more than one stream for the same snippet allows a snippet to be sent in more than one form. For example, the LLT 440 may send a snippet in its actual form through one stream and in a form with bytes complemented and in reverse order through the other stream. The alternating use of different transformations of a snippet more evenly distributes transmission errors among the bits of the snippet as they are received, and hence facilitates the reconstruction of a snippet from multiple corrupted received versions. The receiver always knows which form of the snippet was transmitted based on its stream number.

The LLT 440 partitions the streams into two disjoint subsets, one for use with Hub 110 assigned MAC addresses 750 and the other for use with attaching PEAs' self-selected MAC addresses (AMACs) 760. Both the LLT 440 and the LLD 450 know the size and direction of each stream, but the LLT 450 is responsible for determining how the streams are used, how MAC numbers are assigned and used, and assuring that no two PEAs 120 respond to the same token (containing a MAC address and stream number) transmitted by the Hub 110. One exception to

this includes the Hub's use of its MAC address to broadcast its heartbeat 770 (described below) to all PEAs 120.

EXEMPLARY COMMUNICATION

Fig. 8 is an exemplary diagram of a TDMA frame structure 800 of a TDMA plan consistent with the present invention. The TDMA frame 800 starts with a beacon 810, and then alternates token broadcasts 820 and data transfers 830. The Hub 110 broadcasts the beacon 810 at the start of each TDMA frame 800. The PEAs 120 use the beacon 810, which may contain a unique identifier of the Hub 110, to synchronize to the Hub 110.

Each token 640 (Fig. 6) transmitted by the Hub 110 in a token broadcast 820 includes a MAC address 610 (Fig. 6) and a stream number 620 for the data buffer 630 transfer that follows.

The MAC address 610 and stream number 620 in the token 640 together specify a particular PEA 120 to transmit or receive data, or, in the case of the Hub's MAC address 610, specify no, many, or all PEAs to receive data from the Hub 110 (depending on the stream number). The stream number 620 in the token 640 indicates the direction of the data transfer 830 (Hub 110 to PEA 120 or PEA 120 to Hub 110), the number of bytes to be transferred, and the data source (for the sender) and the appropriate empty data block (for the receiver).

The TDMA plan controls the maximum number of bytes that can be sent in a data transfer 830. Not all of the permitted bytes need to be used in the data transfer 830, however, so the Hub 110 may schedule a status block in the initial segment of a TDMA time interval that is large enough to send a snippet. The Hub 110 and PEA 120 treat any left over bytes as no-ops to mark time. Any PEA 120 not involved in the data transfer uses all of the data transfer 830 bytes

to mark time while waiting for the next token 640. The PEA 120 may also power down non-essential circuitry at this time to reduce power consumption.

Fig. 9A is an exemplary diagram of communication processing for transmitting a single data block from the Hub 110 to a PEA 120 according to the TDMA plan of Fig. 8. Figs. 9B and 9C are flowcharts of the Hub 110 and PEA 120 activities, respectively, of Fig. 9A. The reference numbers in Fig. 9A correspond to the flowchart steps of Figs. 9B and 9C.

With regard to the Hub activity, the Hub 110 responds to a token command in the TDMA plan [step 911] (Fig. 9B) by determining the location of the next data block 600 to send or receive [step 912]. The Hub 110 reads the block's MAC address 610 and stream number 620 [step 913] and generates a token 640 from the MAC address and stream number using FEC [step 914]. The Hub 110 then waits for the time for sending a token 640 in the TDMA plan (i.e., a token broadcast 820 in Fig. 8) [step 915] and broadcasts the token 640 to the PEAs 120 [step 916]. If the stream number 620 in the token 640 is zero (i.e., a NO-DATA-TRANSFER token), no PEA 120 will respond and the Hub 110 waits for the next token command in the TDMA plan [step 911].

If the stream number 620 is non-zero, however, the Hub 110 determines the size and direction of the data transmission from the stream number 620 and waits for the time for sending the data in the TDMA plan (i.e., a data transfer 830) [step 917]. Later, when instructed to do so by the TDMA plan (i.e., after the PEA 120 identified by the MAC address 610 has had enough time to prepare), the Hub 110 transmits the contents of the data buffer 630 [step 918]. The Hub 110 then prepares for the next token command in the TDMA plan [step 919].

With regard to the PEA activity, the PEA 120 reaches a token command in the TDMA plan [step 921] (Fig. 9C). The PEA 120 then listens for the forward error-corrected token 640, having a MAC address 610 and stream number 620, transmitted by the Hub 110 [step 922]. The PEA 120 decodes the MAC address from the forward error-corrected token [step 923] and, if it is not the PEA's 120 MAC address, sleeps through the next data transfer 830 in the TDMA plan [step 924]. Otherwise, the PEA 120 also decodes the stream number 620 from the token 640.

All PEAs 120 listen for the Hub heartbeat that the Hub 110 broadcasts with a token containing the Hub's MAC address 610 and the heartbeat stream 770. During attachment (described in more detail below), the PEA 120 may have two additional active MAC addresses 610, the one it selected for attachment and the one the Hub 110 assigned to the PEA 120. The streams are partitioned between these three classes of MAC addresses 610, so the PEA 120 may occasionally find that the token 640 contains a MAC address 610 that the PEA 120 supports, but that the stream number 620 in the token 640 is not one that the PEA 120 supports for this MAC address 610. In this case, the PEA 120 sleeps through the next data transfer 830 in the TDMA plan [step 924].

Since the PEA 120 supports more than one MAC address 610, the PEA 120 uses the MAC address 610 and the stream number 620 to identify a suitable empty data block [step 925]. The PEA 120 writes the MAC address 610 and stream number 620 it received in the token 640 from the Hub 110 into the data block [step 926]. The PEA 120 then determines the size and direction of the data transmission from the stream number 620 and waits for the transmission of the data buffer 630 contents from the Hub 110 during the next data transfer 830 in the TDMA

plan [step 927]. The PEA 120 stores the data in the data block [step 928], and then prepares for the next token command in the TDMA plan [step 929].

Figs. 9A-9C illustrate communication of a data block from the Hub 110 to a PEA 120.

When the PEA 120 transfers a data block to the Hub 110, similar steps occur except that the Hub 110 first determines the next data block to receive (with its MAC address 610 and stream number 620) and the transmission of the data buffer 630 contents occurs in the opposite direction. The Hub 110 needs to arrange in advance for receiving data from PEAs 120 by populating the MAC address 610 and stream number 620 into data blocks with empty data buffers 630, because the Hub 110 generates the tokens for receiving data as well as for transmitting data.

Figs. 10A and 10B are high-level diagrams of the states that the Hub 110 and PEA 120 LLT 440 (Fig. 4) go through during a data transfer in an implementation consistent with the present invention. Fig. 10A illustrates states of a Hub-to-PEA transfer and Fig. 10B illustrates states of a PEA-to-Hub transfer.

During the Hub-to-PEA transfer (Fig. 10A), the Hub 110 cycles through four states: fill, send even parity, fill, and send odd parity. The fill states indicate when the NI 430 (Fig. 4) may fill a data snippet. The even and odd send states indicate when the Hub 110 sends even numbered and odd numbered snippets to the PEA 120. The PEA 120 cycles through two states: want even and want odd. The two states indicate the PEA's 120 desire for data, with 'want even' indicating that the last snippet successfully received had odd parity. The PEA 120 communicates its current state to the Hub 110 via its status messages (i.e., the state changes serve as ACKs). The Hub 110 waits for a state change in the PEA 120 before it transitions to its next fill state.

During the PEA-to-Hub transfer (Fig. 10B), the Hub 110 cycles through six states: wait/listen for PEA-ready-to-send-even status, read even, send ACK and listen for status, wait/listen for PEA-ready-to-send-odd status, read odd, and send ACK and listen for status. According to this transfer, the PEA 120 cannot transmit data until the Hub 110 requests data, which it will only do if it sees from the PEA's status that the PEA 120 has the next data block ready.

The four listen for status states schedule when the Hub 110 asks to receive a status message from the PEA 120. The two 'send ACK and listen for status' states occur after successful receipt of a data block by the Hub 110, and in these two states the Hub 110 schedules both the sending of Hub status to the PEA 120 and receipt of the PEA status. The PEA status informs the Hub 110 when the PEA 120 has successfully received the Hub 110 status and has transitioned to the next 'fill' state.

Once the PEA 120 has prepared its next snippet, it changes its status to 'have even' or 'have odd' as appropriate. When the Hub 110 detects that the PEA 120 has advanced to the fill state or to 'have even/odd,' it stops scheduling the sending of Hub status (ACK) to the PEA 120. If the Hub 110 detects that the PEA 120 is in the 'fill' state, it transitions to the following 'listen for status' state. If the PEA 120 has already prepared a new snippet for transmission by the time the Hub 110 learns that its ACK was understood by the PEA 120, the Hub 110 skips the 'listen for status' state and moves immediately to the next appropriate 'read even/odd' state. In this state, the Hub 110 receives the snippet from the PEA 120.

The PEA 120 cycles through four states: fill, have even, fill, and have odd (i.e., the same four states the Hub 110 cycles through when sending snippets). The fill states indicate when the

NI 430 (Fig. 4) can fill a data snippet. During the fill states, the PEA 110 sets its status to 'have nothing to send.' The PEA 120 does not transition its status to 'have even' or 'have odd' until the next snippet is filled and ready to send to the Hub 110. These two status states indicate the parity of the snippet that the PEA 120 is ready to send to the Hub 110. When the Hub 110 receives a status of 'have even' or 'have odd' and the last snippet it successfully received had the opposite parity, it schedules the receipt of data, which it thereafter acknowledges with a change of status that it sends to the PEA 120.

EXEMPLARY ATTACHMENT PROCESSING

The Hub 110 communicates with only attached PEAs 120 that have an assigned MAC address 610. An unattached PEA can attach to the Hub 110 when the Hub 110 gives it an opportunity to do so. Periodically, the Hub 110 schedules attachment opportunities for unattached PEAs that wish to attach to the Hub 110, using a small set of attach MAC (AMAC) addresses and a small set of streams dedicated to this purpose.

After selecting one of the designated AMAC addresses 610 at random to identify itself and preparing to send a small, possibly forward error-corrected, "attach-interest" message and a longer, possibly checksummed, "attach-request" message using this AMAC and the proper attach stream numbers 620, the PEA 120 waits for the Hub 110 to successfully read the attach-interest and then the attach-request messages. Reading of a valid attach-interest message by the Hub 110 causes the Hub 110 believe that there is a PEA 120 ready to send the longer (and hence more likely corrupted) attach-request.

Once a valid attach-interest is received, the Hub 110 schedules frequent receipt of the attach-request until it determines the contents of the attach-request, either by receiving the block

intact with a valid checksum or by reconstructing the sent attach-request from two or more received instances of the sent attach-request. The Hub 110 then assigns a MAC address to the PEA 120, sending the address to the PEA 120 using its AMAC address.

The Hub 110 confirms receipt of the MAC address by scheduling the reading of a small, possibly forward error-corrected, attach-confirmation from the PEA 120 at its new MAC address 610. The Hub 110 follows this by sending a small, possibly forward error-corrected, confirmation to the PEA 120 at its MAC address so that the PEA 120 knows it is attached. The PEA 120 returns a final small, possibly forward error-corrected, confirmation acknowledgement to the Hub 110 so that the Hub 110, which is in control of all scheduled activity, has full knowledge of the state of the PEA 120. This MAC address remains assigned to that PEA 120 for the duration of the time that the PEA 120 is attached.

Figs. 11 and 12 are flowcharts of Hub and PEA attachment processing, respectively, consistent with the present invention. When the Hub 110 establishes the network, its logic initializes the attachment process and, as long as the Hub 110 continues to function, periodically performs attachment processing. The Hub 110 periodically broadcasts heartbeats containing a Hub identifier (selecting a new heartbeat identifier value each time it reboots) and an indicator of the range of AMACs that can be selected from for the following attach opportunity [step 1110] (Fig. 11). The Hub 110 schedules an attach-interest via a token that schedules a small PEA-to-Hub transmission for each of the designated AMACs, so unattached PEAs may request attachment.

Each attaching PEA 120 selects a new AMAC at random from the indicated range when it hears the heartbeat. Because the Hub 110 may receive a garbled transmission whenever more

than one PEA 120 transmits, the Hub 110 occasionally indicates a large AMAC range (especially after rebooting) so that at least one of a number of PEAs 120 may select a unique AMAC 610 and become attached. When no PEAs 120 have attached for some period of time, however, the Hub 110 may select a small range of AMACs 610 to reduce attachment overhead, assuming that

5 PEAs 120 will arrive in its vicinity in at most small groups. The Hub 110 then listens for a valid attach-interest from an unattached PEA [step 1120]. The attach-interest is a PEA-to-Hub message having the AMAC address 610 selected by the unattached PEA 120.

Upon receiving a valid attach interest, the Hub 110 schedules a PEA-to-Hub attach-request token with the PEA's AMAC 610 and reads the PEA's attach-request [step 1130]. Due

10 to the low-power wireless environment of the PAN 100, the attach-request transmission may take more than one attempt and hence may require scheduling the PEA-to-Hub attach-request token more than once. When the Hub 110 successfully receives the attach-request from the PEA, it assigns a MAC address to the PEA [step 1140]. In some cases, the Hub 110 chooses the MAC address from the set of AMAC addresses.

15 The Hub 110 sends the new MAC address 610 in an attach-assignment message to the now-identified PEA 120, still using the PEA's AMAC address 610 and a stream number 620 reserved for this purpose. The Hub 110 schedules and listens for an attach-confirmation response from the PEA 120 using the newly assigned MAC address 610 [step 1150].

Upon receiving the confirmation from the PEA 120, the Hub 110 sends its own

20 confirmation, acknowledging that the PEA 120 has switched to its new MAC, to the PEA 120 and waits for a final acknowledgment from the PEA 120 [step 1160]. The Hub 110 continues to send the confirmation until it receives the acknowledgment from the PEA 120 or until it times

out. In each of the steps above, the Hub 110 counts the number of attempts it makes to send or receive, and aborts the attachment effort if a predefined maximum number of attempts is exceeded. Upon receiving the final acknowledgment, the Hub 110 stops sending its attach confirmation, informs its NI 430 (Fig. 4) that the PEA 120 is attached, and begins exchanging
5 both data and keep-alive messages (described below) with the PEA 120.

When an unattached PEA 120 enters the network, its LLC 420 (Fig. 4) instructs its LLT 440 to initialize attachment. Unlike the Hub 110, the PEA 120 waits to be polled. The PEA 120 instructs its DCL 460 to activate and associate the heartbeat stream 770 (Fig. 7B) with the Hub's MAC address and waits for the heartbeat broadcast from the Hub 110 [step 1210] (Fig. 12). The
10 PEA 120 then selects a random AMAC address from the range indicated in the heartbeat to identify itself to the Hub 110 [step 1220]. The PEA 120 instructs its DCL 460 to send an attach-interest and an attach-request data block to the Hub 110, and activate and associate the streams with its AMAC address [step 1230]. The PEA 120 tells its driver to activate and respond to the selected AMAC address for the attach-assignment stream.

15 The unattached PEA 120 then waits for an attach-assignment with an assigned MAC address from the Hub 110 [step 1240]. Upon receiving the attach-assignment, the PEA 120 finds its Hub-assigned MAC address and tells its driver to use this MAC address to send an attach-confirmation to the Hub 110 to acknowledge receipt of its new MAC address [step 1250], activate all attached-PEA streams for its new MAC address, and deactivate the streams
20 associated with its AMAC address.

The PEA 120 waits for an attach confirmation from the Hub 110 using the new MAC address [step 1260] and, upon receiving it, sends a final acknowledgment to the Hub 110 [step 1270]. The PEA 120 then tells its NI 430 that it is attached.

The PEA 120, if it hears another heartbeat from the Hub 110 before it completes attachment, discards any prior communication and begins its attachment processing over again with a new AMAC.

EXEMPLARY DETACHMENT AND REATTACHMENT PROCESSING

The Hub 110 periodically informs all attached PEAs 120 that they are attached by sending them 'keep-alive' messages. The Hub 110 may send the messages at least as often as it transmits heartbeats. The Hub 110 may send individual small, possibly forward error-corrected, keep-alive messages to each attached PEA 120 when few PEAs 120 are attached, or may send larger, possibly forward error-corrected, keep-alive messages to groups of PEAs 120.

Whenever the Hub 110 schedules tokens for PEA-to-Hub communications, it sets a counter to zero. The counter resets to zero each time the Hub 110 successfully receives a block (either uncorrupted or reconstructed) from the PEA 120, and increments for unreadable blocks. If the counter exceeds a predefined threshold, the Hub 110 automatically detaches the PEA 120 without any negotiation with the PEA 120. After this happens, the Hub 110 no longer schedules data or status transfers to or from the PEA 120, and no longer sends it any keep-alive messages.

Fig. 13 is a flowchart of PEA detachment and reattachment processing consistent with the present invention. Each attached PEA 120 listens for Hub heartbeat and keep-alive messages [step 1310]. When the PEA 120 first attaches, and after receiving each keep-alive message, it resets its heartbeat counter to zero [step 1320]. Each time the PEA 120 hears a heartbeat, it

increments the heartbeat counter [step 1330]. If the heartbeat counter exceeds a predefined threshold, the PEA 120 automatically assumes that the Hub 110 has detached it from the network 100 [step 1340]. After this happens, the PEA 120 attempts to reattach to the Hub 110 [step 1350], using attachment processing similar to that described with respect to Figs. 11 and 12.

5 If the Hub 110 had not actually detached the PEA 120, then the attempt to reattach causes the Hub 110 to detach the PEA 120 so that the attempt to reattach can succeed. When the PEA 120 is out of range of the Hub 110, it may not hear from the Hub 110 and, therefore, does not change state or increment its heartbeat counter. The PEA 120 has no way to determine whether the Hub 110 has detached it or how long the Hub 110 might wait before detaching it. When the
10 PEA 120 comes back into range of the Hub 110 and hears the Hub heartbeat (and keep-alive if sent), the PEA 120 then determines whether it is attached and attempts to reattach if necessary.

CONCLUSION

Systems and methods consistent with the present invention provide a wireless personal area network that permit a host device to communicate with a varying number of peripheral
15 devices with minimal power and minimal interference from neighboring networks by using a customized TDMA protocol. The host device uses tokens to facilitate the transmission of data blocks through the network.

The foregoing description of exemplary embodiments of the present invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the
20 precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention. The scope of the invention is defined by the claims and their equivalents.

WHAT IS CLAIMED IS:

1. A network comprising:
a hub device configured to generate a token and broadcast the token on the network; and
at least one peripheral device configured to receive the token broadcast by the hub device,
determine whether the token identifies the peripheral device, analyze the token to determine a
5 size and direction of a current data transfer when the token identifies the peripheral device, and
transfer data to or receive data from the hub device according to the determined size and
direction of the current data transfer.

2. The network of claim 1, further comprising:
a single wireless communication channel having a plurality of logical unidirectional
communication streams, the data transfer occurring over one of the communication streams.

3. The network of claim 2, wherein the token includes:
an address of one of the hub device and the peripheral device, and
a stream number that identifies one of the communication streams.

4. The network of claim 2, wherein each of communication streams has a
predetermined size and direction of a data transfer.

5. The network of claim 1, wherein the network operates according to a
communications protocol shared by the hub device and the peripheral device to synchronize

timing of communications.

6. The network of claim 5, wherein the communications protocol includes a plurality of frames, each of the frames including:

a signal that marks a start of the frame,

at least one token transmission that identifies the peripheral device, and

5 at least one data transfer opportunity that permits the hub device to communicate a data block with the peripheral device.

7. The network of claim 1, wherein the hub device includes:

a plurality of data blocks, each of the data blocks containing an address of one of the hub device and a peripheral device, a stream number that specifies a size and direction of transfer, and a data buffer, the hub device cycling through the data blocks to generate the tokens.

8. The network of claim 1, wherein the peripheral device is further configured to respond to at least two addresses.

9. The network of claim 8, wherein the peripheral device is further configured to associate at least one active communication stream with each of the at least two addresses.

10. The network of claim 1, wherein the peripheral device is further configured to respond to at least three addresses, one of the three addresses being an address of the hub device.

11. The network of claim 1, wherein the peripheral device includes a plurality of virtual peripheral devices.

12. The network of claim 1, wherein the hub device includes a memory including:
a link layer control structure that performs network bandwidth control and token planning,

a network interface structure that determines whether and when to schedule a data transfer, and

a link layer transport structure that provides a reliable data transfer for the network interface.

13. The network of claim 1, wherein the peripheral device includes a memory including:

a link layer control structure that performs token planning,

a network interface structure that determines whether and when to schedule a data

transfer, and

a link layer transport structure that provides a reliable data transfer for the network interface.

14. The network of claim 1, wherein at least one of the hub device and the peripheral device is further configured to transfer data in multiple forms.

15. The network of claim 14, wherein at least one of the hub device and the peripheral device is further configured to transmit the data in an original form and at least one of a complemented form and a reverse order form.

16. The network of claim 14, wherein at least one of the hub device and the peripheral device is further configured to differentiate the multiple forms of the data by a communications stream on which the data was transmitted.

17. The network of claim 14, wherein at least one of the hub device and the peripheral device is further configured to combine the multiple forms of the data to reconstruct the data.

18. The network of claim 1, wherein the hub device is further configured to schedule transmission of a status block from the peripheral device.

19. The network of claim 18, wherein the hub device is further configured to schedule transmission of data from the peripheral device when the status block from the peripheral device indicates that the peripheral device has data ready for transmission to the hub device.

20. A network comprising:

hub means for generating a token and for broadcasting the token on the network; and

peripheral means connected to the hub means, including

means for receiving the token broadcast by the hub means,

5 means for determining whether the token identifies the peripheral means,
means for analyzing the token to determine a size and direction of a current data
transfer when the token identifies the peripheral means, and
means for transferring data to or receiving data from the hub means according to
the determined size and direction of the current data transfer.

10

21. A method for transmitting data in a network having a hub device connected to at
least one peripheral device, comprising:

generating a token;

broadcasting, by the hub device, the token on the network;

5 receiving the broadcast token at the peripheral device;

determining, at the peripheral device, that the token identifies the peripheral device;

analyzing the token to determine a size and direction of a current data transfer when the
token identifies the peripheral device; and

transferring data between the peripheral device and the hub device according to the
10 determined size and direction of the current data transfer.

22. The method of claim 21, wherein the network includes a plurality of wireless
unidirectional communication streams; and

wherein the transferring data includes:

sending data over one of the communication streams.

23. The method of claim 21, wherein the generating a token includes:
accessing a data block in the hub device to identify an address and a communication stream for the current data transfer, and
generating the token based on the identified address and communication stream.
24. The method of claim 23, wherein the determining includes:
decoding the token to identify the address and the communication stream, and
analyzing the identified address to determine whether the identified address matches an address of the peripheral device.
25. The method of claim 24, wherein the analyzing the identified address includes:
determining whether the identified communication stream is active for the identified address.
26. The method of claim 21, wherein the transferring includes:
transmitting the data in multiple forms.
27. The method of claim 26, wherein the transmitting includes:
sending the data in an original form and at least one of a complemented form and a reverse order form.

28. The method of claim 26, further comprising:
differentiating the multiple forms of the data by a communications stream on which the data was transmitted.

29. The method of claim 26, further comprising:
combining the multiple forms of the data to reconstruct the data.

30. The method of claim 21, wherein the transferring includes:
scheduling, by the hub device, transmission of a status block from the peripheral device.

31. The method of claim 30, wherein the transferring further includes:
scheduling, by the hub device, transmission of data from the peripheral device when the status block from the peripheral device indicates that the peripheral device has data ready for transmission to the hub device.

32. A memory used by at least one of a hub device and a peripheral device to transfer data between the hub device and a plurality of the peripheral devices over a wireless network, comprising:

a network interface structure configured to determine whether and when to schedule a

5 data transfer;

a link layer transport structure configured to receive data from the network interface structure and arrange for a reliable transmission of the data; and

a link layer driver structure configured to receive the data from the link layer transport structure and transmit the data under control of the link layer transport structure.

10

33. The memory of claim 32, wherein the link layer transport structure is further configured to generate a header having error checking information for the data transfer.

34. The memory of claim 32, wherein the link layer transport structure is further configured to use reverse-direction status transmissions for the data transfer.

35. The memory of claim 34, wherein the link layer transport structure is further configured to use retransmissions for the status transmissions and the data transfers.

36. A method for transferring data in a network connecting a hub device to a set of peripheral devices, the network operating according to a communications protocol having a plurality of alternating token slots and data transfer slots, the method, performed by the hub device, comprising:

5 identifying an address and a communication stream for a current data transfer, the address identifying one of the peripheral devices;

generating a token based on the identified address and communication stream;

broadcasting the token on the network during one of the token slots; and

communicating the data between the identified peripheral device and the hub device on

10 the identified communication stream during one of the data transfer slots.

37. A hub device that communicates data in a network connecting the hub device to a plurality of peripheral devices, the network operating according to a communications protocol having a plurality of alternating token slots and data transfer slots, the hub device comprising:

a memory having instructions for:

5 identifying an address and a communication stream for a current data transfer, the address identifying one of the peripheral devices,

generating a token based on the identified address and communication stream,

broadcasting the token on the network during one of the token slots, and

10 communicating the data between the identified peripheral device and the hub device on the identified communication stream during one of the data transfer slots; and
a processor that executes the instructions in the memory.

38. A method for transferring data in a network connecting at least one peripheral device to a hub device, the method, performed by one of the peripheral devices, comprising:

receiving a token from the hub device that identifies the peripheral device;

analyzing the token to determine a size and direction of a current data transfer; and

5 transferring data to or receiving data from the hub device according to the determined size and direction of the current data transfer.

39. The method of claim 38, wherein the token includes an address and a communication stream identifier; and

wherein the analyzing the token includes:

decoding the token to identify the address and the communication stream
5 identifier,
determining whether the address identifies the peripheral device, and
determining the size and direction of the current data transfer from the
communication stream identifier.

40. A peripheral device that communicates data in a network connecting at least one
peripheral device to a hub device, the peripheral device comprising:

a memory that stores instructions; and
a processor that executes the instructions in the memory to receive a token from the hub
5 device that identifies the peripheral device, analyze the token to determine a size and direction of
a current data transfer, and transfer data to or receive data from the hub device according to the
determined size and direction of the current data transfer.

41. A communications protocol used in a network connecting a hub device to at least
one peripheral device, the communications protocol having a plurality of frames comprising:

a beacon that marks a start of one of the frames;
at least one token transmission that identifies one of the peripheral devices for a data

5 transfer; and

at least one data transfer opportunity that permits the hub device to communicate a data
block with the identified peripheral device.

42. The communications protocol of claim 41, wherein the token transmission includes an address of the identified peripheral device and a communication stream identifier that identifies a communication stream for the data transfer.

43. The communications protocol of claim 42, wherein the communication stream identifier includes data regarding a size and direction of the data transfer.

ABSTRACT

A network (100) connects a hub device (110) to several peripheral devices (120). The hub device (110) generates a token and broadcasts the token on the network (100). Each of the peripheral devices (120) receives the token broadcast by the hub device (110), determines

- 5 whether the token identifies the peripheral device (120), analyzes the token to determine a size and direction of a current data transfer when the token identifies the peripheral device (120), and transfers data to or receives data from the hub device (110) according to the determined size and direction of the current data transfer.

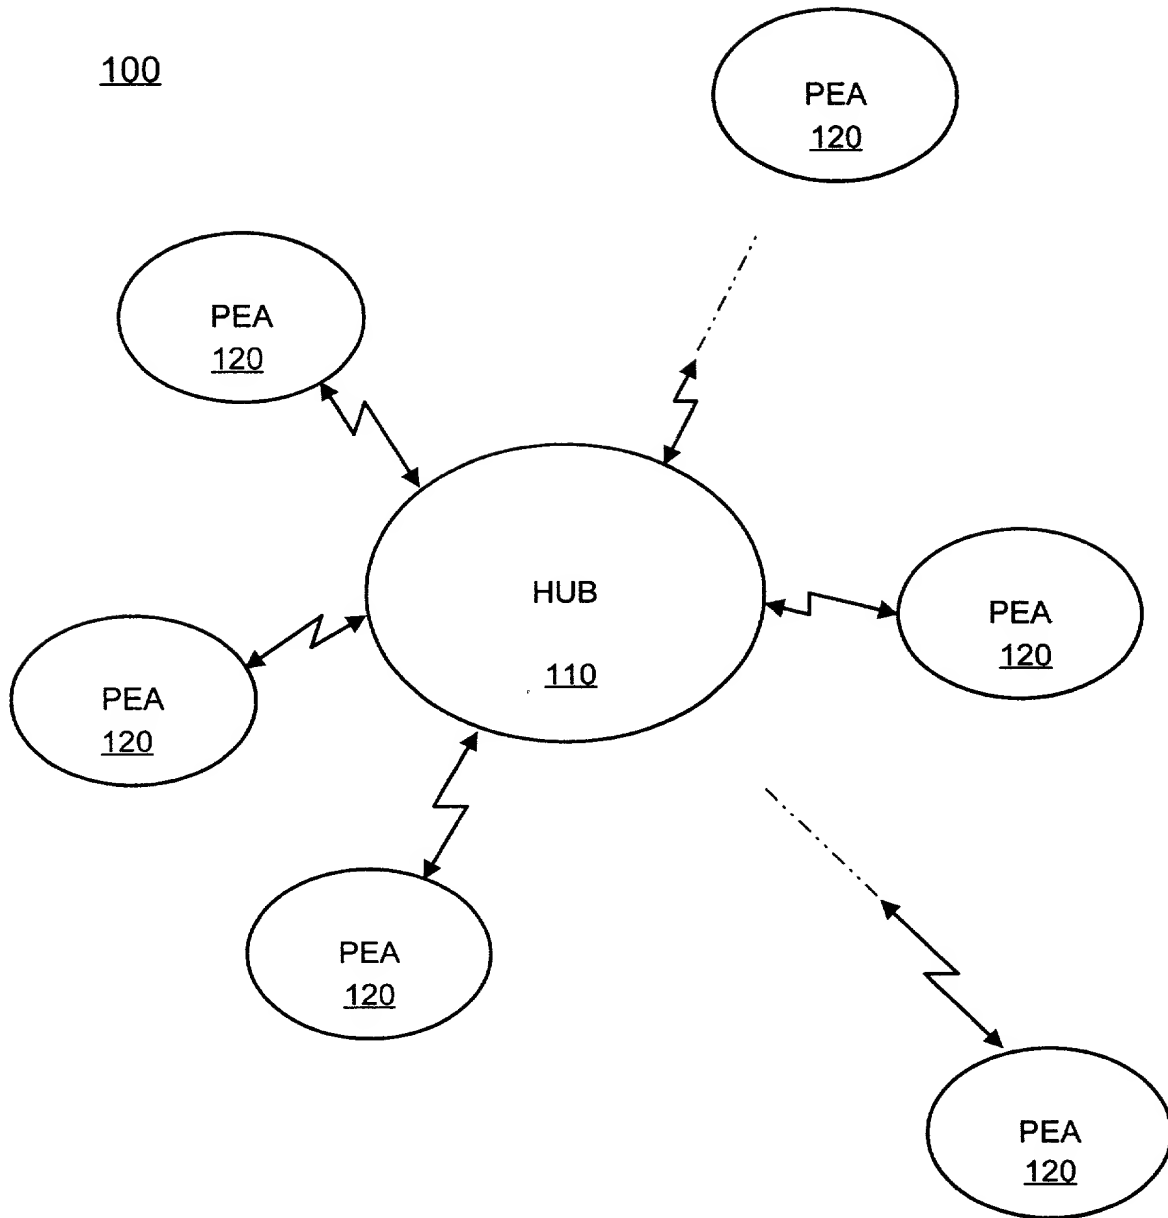


FIG. 1




FIG. 2

120




FIG. 3




FIG. 4

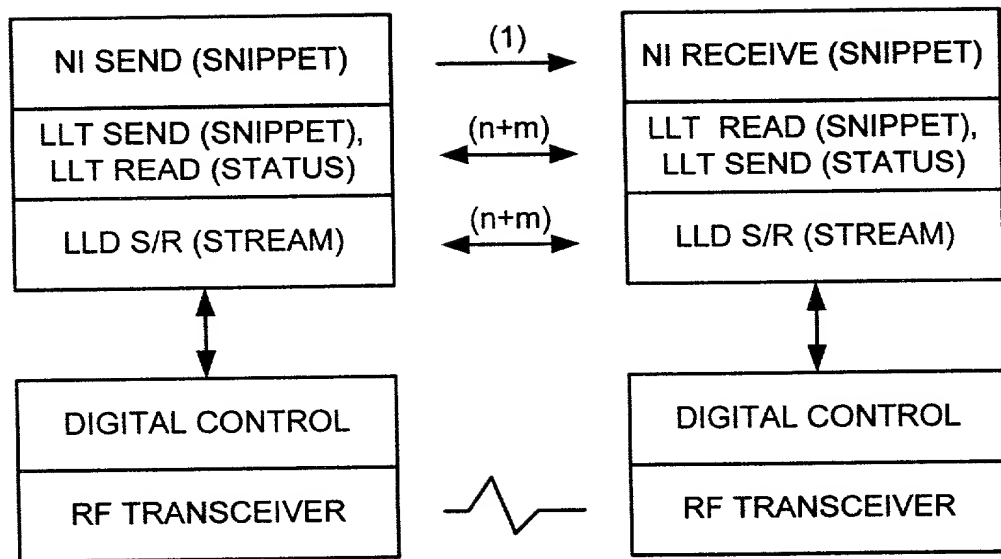


FIG. 5




FIG. 6

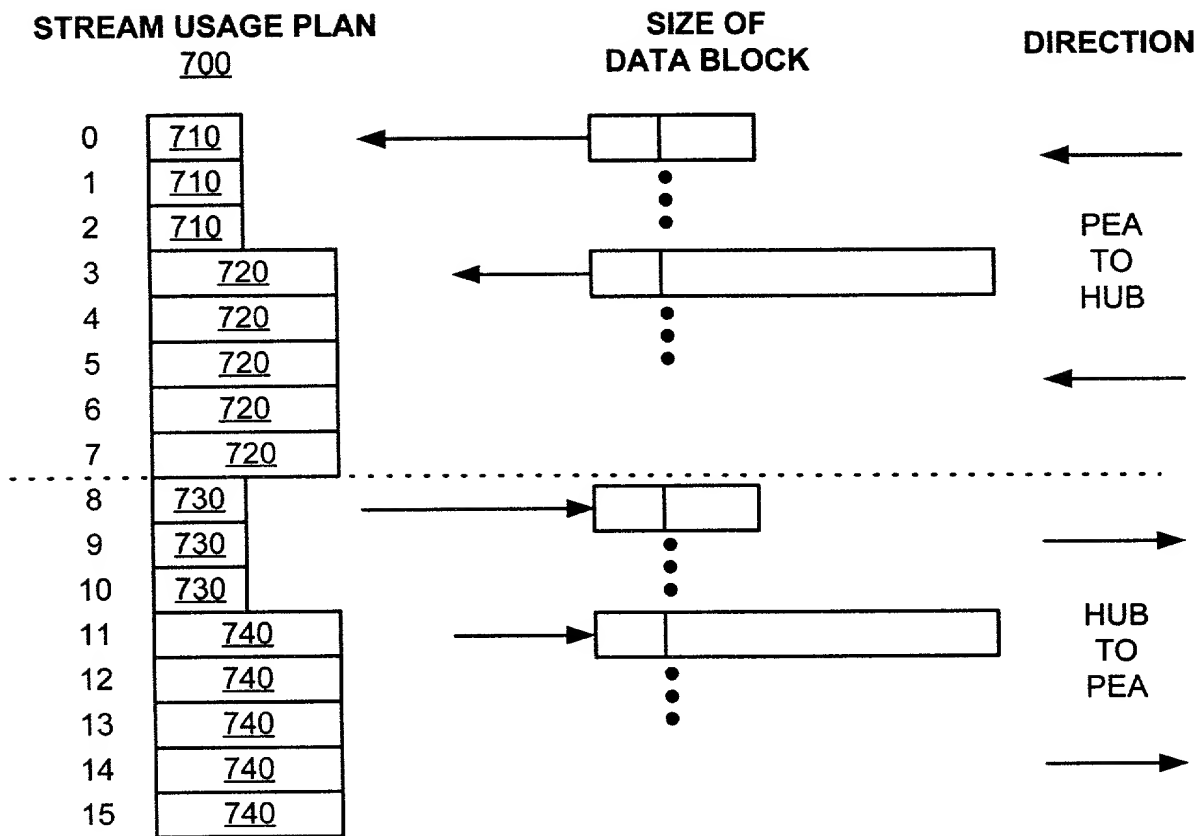


FIG. 7A

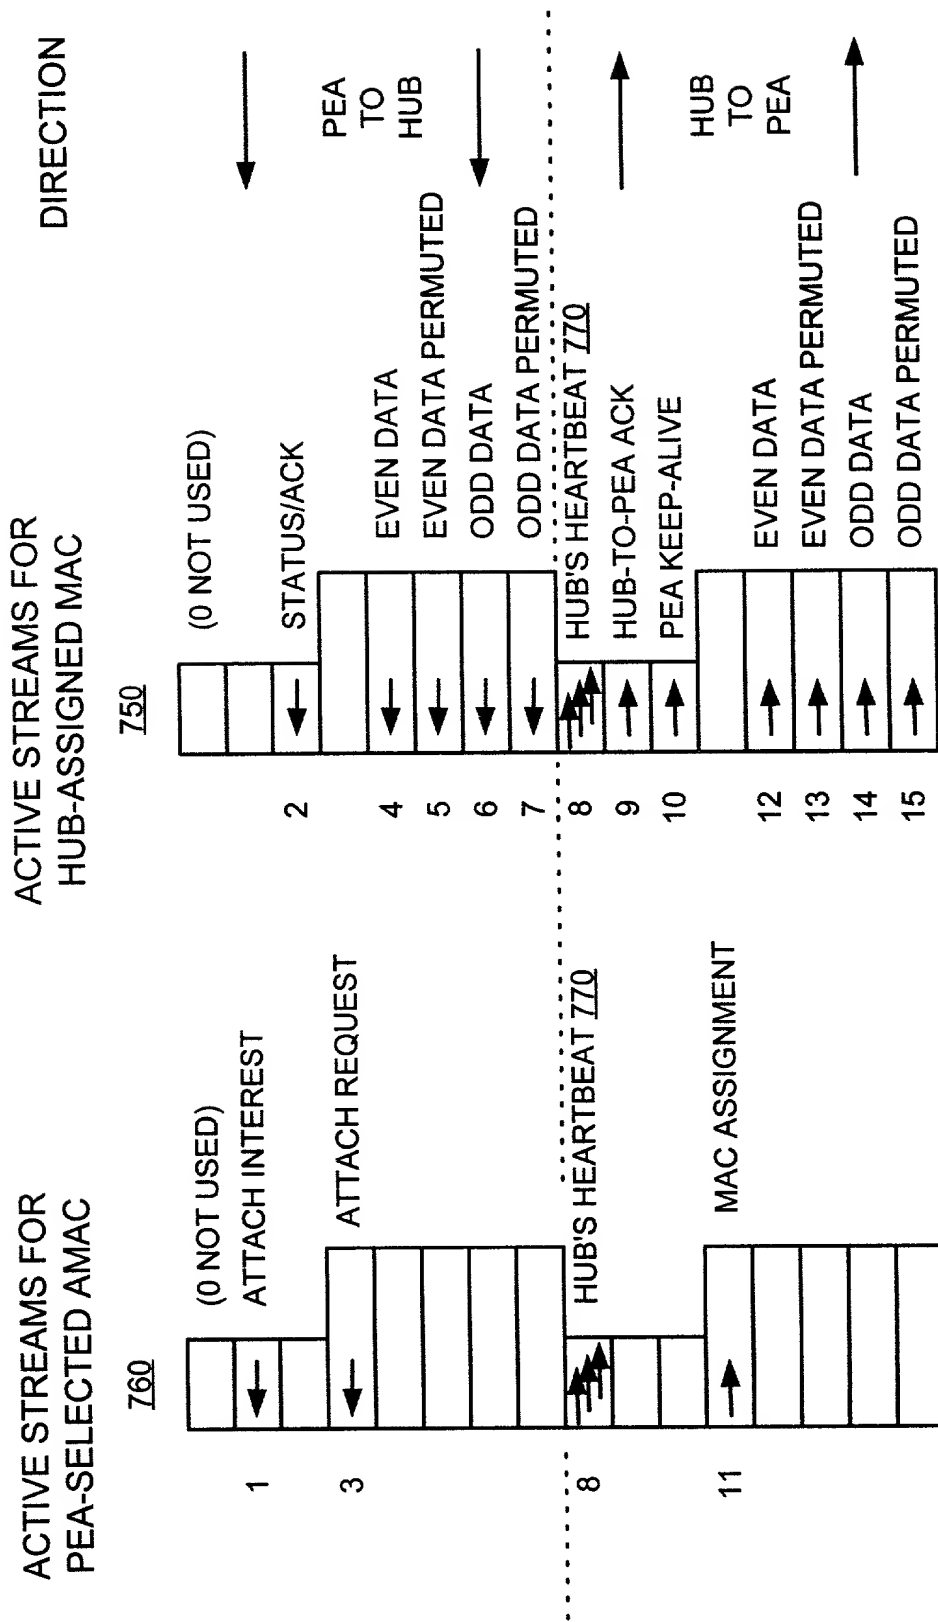


FIG. 7B

800

TDMA FRAME STRUCTURE

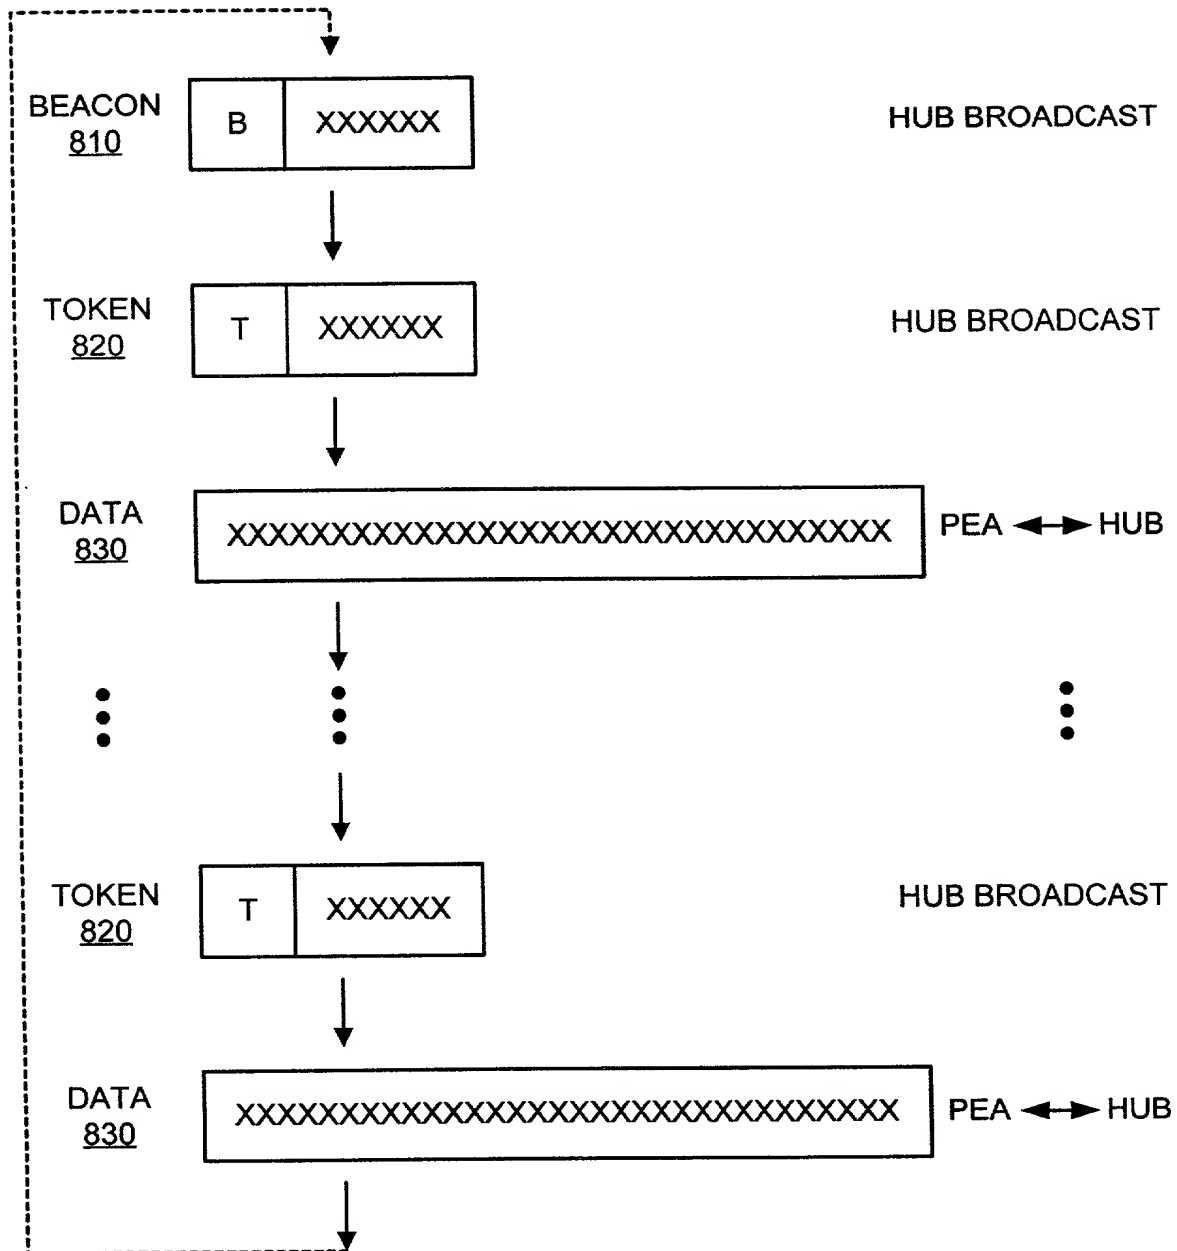


FIG. 8





FIG. 9A






FIG. 9C

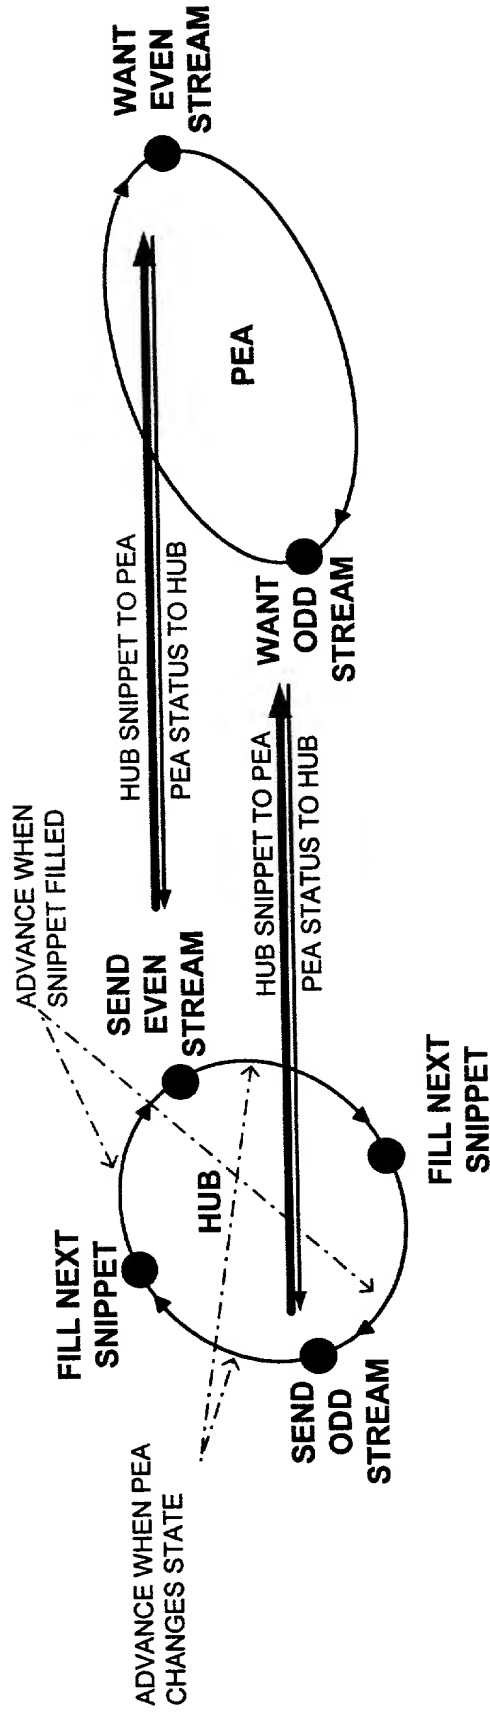


FIG. 10A






FIG. 12

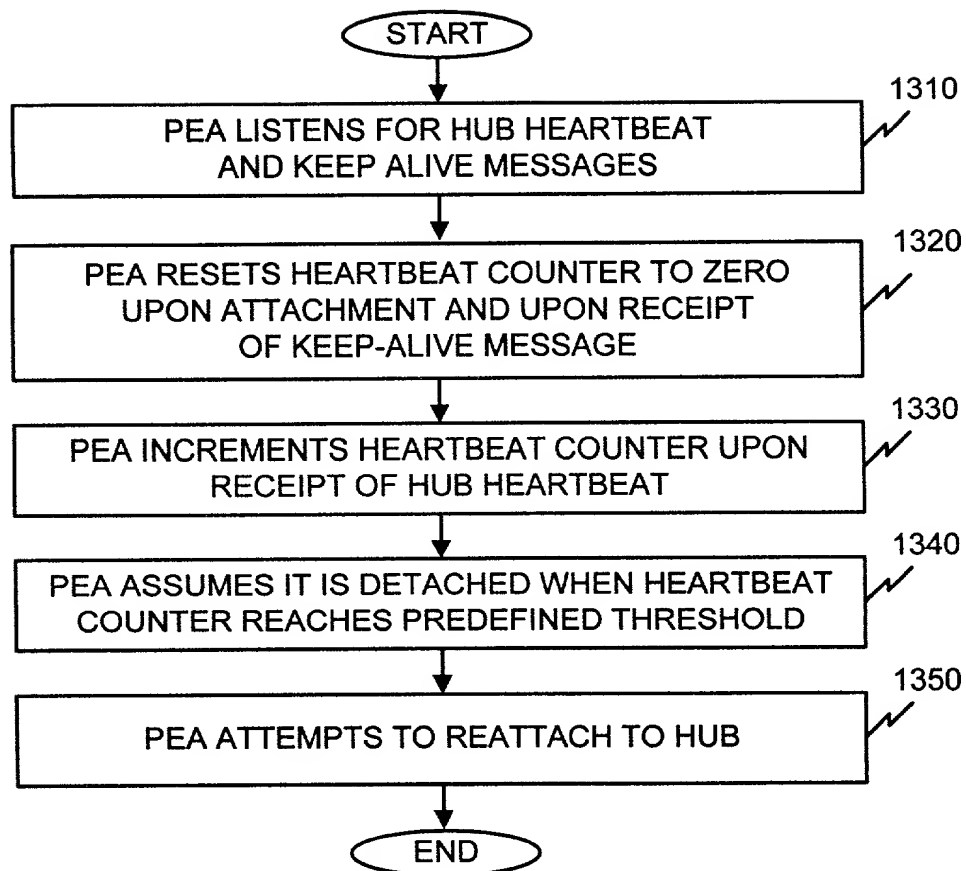


FIG. 13

**DECLARATION AND POWER OF ATTORNEY FOR PATENT
APPLICATION**

Attorney Docket No. 99-412

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

SYSTEMS AND METHODS FOR COMMUNICATING IN A PERSONAL AREA NETWORK

the specification of which (check one) ☒ is attached hereto. ☐ was filed on

as Appln. Serial No.

and was amended on

(if applicable). I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Claimed

☐ Yes ☐ No

(Number)

(Country)

(Day/Month/Year filed)

I hereby claim the benefit under Title 35, United States Code, 119(e) of any United States provisional applications(s) listed below.

(Application Number)

(Filing Date)

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56 which occurred between the filing date of the prior application and the national or PCT international filing date for this application:

(Appln. Serial No.)

(Filing Date)

(Status--patented, pending, abandoned)

Attorney Docket No. 99-412

I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

**Leonard C. Suchyta, Reg. No. 25,707, Floyd E. Anderson, Reg. No. 33,825 and W. Eric Webostad,
Reg. No. 35,406**

Address all telephone calls to Floyd E. Anderson at telephone no. (781) 466-4016

Address all correspondence to Leonard C. Suchyta
GTE Service Corporation
600 Hidden Ridge, HQE03G13
Irving, TX 75038

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR FIRST INVENTOR Robert J. Donaghey

Inventor's signature Robert J. Donaghey Date March 23, 2000

Residence Lexington, Massachusetts Citizenship USA

Post Office Address 40 Oak St., Lexington, MA 02421